

# Hairpin Vortex Formation, a Case Study for Unsteady Visualization

H. M. Tufo,<sup>1</sup> P. F. Fischer,<sup>2</sup> M. E. Papka,<sup>2</sup> M. Szymanski<sup>2</sup>

July 28, 1999

## Abstract

To better understand the vortex dynamics of coherent structures in turbulent and transitional boundary layers, we consider direct numerical simulation of the interaction between a flat-plate-boundary-layer flow and an isolated hemispherical roughness element. Of principal interest is the evolution of hairpin vortices that form an interlacing pattern in the wake of the hemisphere, lift away from the wall, and are stretched by the shearing action of the boundary layer. Using animations of unsteady three-dimensional representations of this flow, produced by the vtk toolkit and enhanced to operate in a CAVE virtual environment, we identify and study several key features in the evolution of this complex vortex topology not previously observed in other visualization formats.

## 1 Introduction

Fluid flow is an inherently visual phenomenon. In our daily experience it holds a special visual fascination, to the point where the term *fluid* itself has become a descriptor of motion and shapes. Early quantitative experiments in fluid mechanics were largely based on or inspired by qualitative visualization. For instance, Osborne Reynolds's famous demonstration of the transition of flow in a pipe from a laminar to a turbulent state was based on the observation that a thin trace of dye introduced into the flow would suddenly expand into a chaotic cloud when the flow speed was increased.

Ironically, this branch of science, so amenable to visual observation and so pervasive in our daily lives, has proven quite difficult to quantify in detail. Advanced experimental methods, such as particle image velocimetry, currently allow temporal quantification of unsteady velocity fields in only two-dimensional slices. Three-dimensional histories of entire flow fields are still beyond the reach of experimental methods. Numerical methods for the solution of the governing Navier-Stokes equations have been under development for almost a century, with adequate computing power for full, unsteady three-dimensional flows becoming available only in the past two decades. While such simulations have permitted detailed pointwise interrogation of a flow field, the amount of data generated by unsteady flows has made it difficult to visualize the dynamic evolution of *flow structures*. However, the advent of powerful three-dimensional graphics systems with closely coupled compute engines now allows interactive visualization of numerically generated

---

<sup>1</sup>Center on Astrophysical Thermonuclear Flashes, University of Chicago, Chicago, IL 60637.

<sup>2</sup>Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439.




Figure 1: Computational domain showing inlet velocity profile, flat plate, hemisphere, and isolated hairpin vortex in hemisphere wake. For clarity, the vortex has been reflected about the symmetry plane.

unsteady three-dimensional flow fields in an immersive environment that provides visual cues similar to those that led to Osborne Reynolds’s insight over a century ago.

In this paper, we present flow visualization results of numerically generated hairpin vortex formation and evolution. At moderate flow speeds, hairpin vortices provide an example of an observable, organized transition process from steady, two-dimensional laminar flows to unsteady, three-dimensional turbulent flows. Consequently, hairpin vortices are of interest in the study of transitional and turbulent boundary layers, where they have been frequently observed experimentally. In this study, the hairpin vortices are generated in the wake of a hemispherical roughness element embedded in a flat-plate boundary layer, following closely the experiments of Acalar and Smith [1] and, to a lesser extent, those of Klebanoff et al. [12]. Numerical studies of hairpin vortices have also been undertaken by Singer under slightly different flow conditions [2, 16, 17].

The basic flow configuration is shown in Fig. 1. A time-independent velocity profile is prescribed across the upstream entrance of the domain. Based on experimental observations [1], the flow is symmetric about the plane  $y=0$ , so the flow is computed only in the half-domain shown. At sufficiently high nondimensional flow speeds, or Reynolds numbers, the steady boundary layer region near the plate is destabilized by the hemisphere, resulting in periodic shedding of hairpin vortices in the wake.

The simulation results are analyzed by using X-window-based software developed specifically for the numerical methods employed and using interactive software built on top of the “Visualization Toolkit” (vtk) to drive stereo visualization environments such as the CAVE<sup>TM</sup> (Cave Automatic Virtual Environment) and ImmersaDesk<sup>TM</sup>. The resulting immersive visualization proceeds at a sufficiently high frame rate that the hairpin evolution can be readily integrated by eye, thereby allowing one a comprehensive understanding of the dynamics of this complex flow.

The remainder of this paper is organized as follows. Section 2 provides a brief overview of the numerical method and specific flow conditions considered. Section 3 describes the immersive visualization environment and software tools used. Section 4 presents a mixture of quantitative and qualitative visual

results used to analyze this flow. We close in Section 5 with a summary of our results.

## 2 Numerical Method and Flow Simulation

The hairpin vortex simulation is based on numerical integration of the unsteady incompressible Navier-Stokes equations,

$$\begin{aligned}\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} &= -\nabla P + \frac{1}{Re} \nabla^2 \mathbf{u} \\ -\nabla \cdot \mathbf{u} &= 0,\end{aligned}$$

coupled with appropriate boundary conditions on the velocity,  $\mathbf{u}$ . Semi-implicit time stepping is employed in which the nonlinear convective terms are treated explicitly, while the viscous and pressure terms are treated implicitly. Spatial discretization is based on the spectral element method, which is a high-order weighted residual technique similar to the finite element method. Within each element, basis functions are based on tensor-products of  $N$ th-order Lagrange polynomials [5, 6, 13]. The nodes of the Lagrange polynomials are taken to be the Gauss-Lobatto-Legendre (GLL) quadrature points, so that high-order GLL quadrature can be substituted for the integrals required for the residual evaluation. This strategy leads to significant sparsity in the resultant matrix operators and, in particular, yields a diagonal mass matrix. For problems having smooth solutions, such as the incompressible Navier-Stokes equations, convergence is exponential in  $N$ , so that little numerical dissipation or dispersion results. This is ideal for transitional flows of the type considered here, which can be very sensitive to such nonphysical effects.

The spatial discretization, coupled with the semi-implicit time stepping, leads to a number of symmetric positive-definite linear systems to be solved at each time step. There is one Helmholtz problem for each velocity component, and a Poisson-like problem for the pressure. These are solved with preconditioned conjugate gradient iteration, using Jacobi preconditioning for the diagonally dominant Helmholtz problems, and an overlapping Schwarz additive procedure developed by Dryja and Widlund for the pressure problem [4, 5, 6]. Work is further reduced by using a high-quality initial approximation to the pressure at each step. This approximation is generated at low cost as a projection onto a few ( $\sim 1$  to 40) previous solutions, and typically reduces the iteration count by a factor of two to four, depending on the problem and iteration tolerances. All of these steps are readily parallelized, save for a “coarse-grid” solve associated with the additive Schwarz procedure. This can become problematic when the number of processors grows much beyond 100, as it involves all-to-all communication with little work per processor. To address this, we have developed a fast parallel coarse-grid solver that involves minimal communication and readily scales to at least 10,000 processors [19].

We considered several meshes and computational domains to test grid and domain independence, with our two primary production meshes having  $K = 1021$  and  $K = 1535$  spectral elements. All of the meshes we tested had a unit-radius hemisphere located with its center at  $(x, y, z) = (0, 0, .1)$ . A short cylinder of height  $dz = .1$  connects the hemisphere to the flat plate, which is located at  $z = 0$ . The cylinder was added so that a thin layer of elements could be placed near the plate to ensure that the boundary layer was adequately resolved. The resultant 10% increase in the height of the hemisphere closely corresponds




Figure 2: (a) Transition from  $Re = 675$  to  $700$ . (b) Strouhal number versus Reynolds number for the  $K = 1021$  and  $K = 1535$  meshes.

to height increases resulting from glue used to fasten the hemispheres in the experiments of Klebanoff et al. [12]. We consider a Reynolds number range of  $Re_R := UR/\nu = 450-700$ , where  $U$  is the free-stream velocity,  $R$  is the hemisphere radius, and  $\nu$  is the kinematic viscosity of the fluid. Flow is in the positive  $x$  direction with incoming flow  $\mathbf{u} = (u_B(y), 0, 0)$ , where  $u_B(y)$  represents a Blasius profile with  $\delta_{99} = 1.2R$ . The inflow boundary is located at  $x = -8.4$  for the  $K = 1021$  mesh and at  $x = -10.0$  for the  $K = 1535$  mesh. To reduce the computational burden, we impose reflection symmetry about the  $y$ -plane. Homogeneous Neumann boundary conditions are prescribed for the velocity at the outlet, located at  $x = 30$ . The upper boundary and the right boundary are also taken to be symmetry planes, corresponding to stress-free boundary conditions. These are located at  $z = -6.5$  and  $y = -6.4$  for  $K = 1021$  and  $z = -8.0$  and  $y = -8.4$  for  $K = 1535$ . We note that, in addition to having different size domains and number of elements, the two meshes have different transitions between spherical and rectangular elements and between the inflow boundary and hemisphere, with the  $K = 1535$  mesh having a smoother transition for both.

To transition between Reynolds numbers, we employ a sine ramp,  $Re(t) = Re_i + (Re_f - Re_i) \sin(gt)$ , where  $Re_i$  is the initial Reynolds number,  $Re_f$  is the final Reynolds number, and  $g$  is the growth rate, typically set such that duration of the ramp is 40–80 convective time units. Figure 2a shows the pressure trace at  $\mathbf{x} = (2.4, 0.0, 1.2)$  for an 80 time unit transition between Reynolds number 675 and 700. Once the transition has been made, the flow is driven until it settles into a steady periodic state. Of particular interest in these simulations is determination of the Strouhal number,  $S = fk/U_k$ , where  $f$  is the shedding frequency,  $k = 1.1$  is the total height of the roughness element, and  $U_k$  is the tip velocity without the hemisphere present.  $U_k$  is calculated from a two-dimensional channel flow simulation on a domain identical in dimension to the hemisphere symmetry plane (centerplane) but without the hemisphere. The frequency,

$f$ , is determined from history plots similar to Fig. 2a. Strouhal numbers for the  $K = 1021$  and 1535 meshes are shown in Fig. 2b. For this study  $3 \leq N \leq 19$ , with the maximum  $N$  considered for Reynolds number 450 being 9 and for Reynolds number 700 being 19. We note that the maximum discrepancy between the two meshes is less than 2% over the range of Reynolds numbers considered here and that spatial and temporal convergence checks were performed at several Reynolds numbers to verify convergence. We further note that our Strouhal numbers compare favorably to those in [1] and [12].

### 3 Scientific Visualization

Two modes of visualization are employed. The first is a menu-driven X-window-based postprocessor developed specifically for the spectral element method. It exploits the full accuracy of the high-order Lagrangian bases when interpolating off-grid point values or computing derived quantities such as velocity gradients or vorticity. In addition, the postprocessor can map spectral element data onto unstructured hexahedral meshes of arbitrary density. This data is then passed into a second visualization package designed for stereo visualization of general mesh data. The primary components of this second mode are the vtk library, which is an open-source software system for visualization that provides a high-level abstraction for constructing scientific visualization applications [15], and the CAVE library, which enables projection and exploration of immersive stereo images [3].

#### 3.1 Vortex Identification

Vortex identification is based on the  $\lambda_2$  method of Jeong and Hussain [9]. Identification of a vortex in viscous flows is challenging because the classic rules governing vortex dynamics generally apply only in the inviscid limit. In boundary layer flows, viscosity is non-negligible, and standard approaches such as integrating vortex lines or using pressure minima or vorticity maxima can lead to improper vortex identification. Jeong and Hussain have established a robust criterion for the identification of vortex (or coherent) structures in viscous flows based on the eigenvalues of the symmetric  $3 \times 3$  tensor

$$A_{ij} := \sum_{k=1}^3 \left( \Omega_{ik} \Omega_{kj} + S_{ik} S_{kj} \right), \quad (1)$$

where

$$\Omega_{ij} := \frac{1}{2} \left( \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \quad S_{ij} := \frac{1}{2} \left( \frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i} \right) \quad (2)$$

represent the symmetric and antisymmetric components of the velocity gradient tensor,  $\nabla \mathbf{u}$ . Given the three (real) eigenvalues of  $A$  at each grid point, a vortex core is identified as any contiguous region having two negative eigenvalues. If the eigenvalues are sorted such that  $\lambda_1 \leq \lambda_2 \leq \lambda_3$ , then any region for which  $\lambda_2 < 0$  corresponds to a vortex core. One advantage of this approach is that vortices can be identified as isosurfaces of a well-defined scalar field. Moreover, the criterion  $\lambda_2(\mathbf{x}) < 0$  is scale invariant, so there is in principle no ambiguity in selecting which isosurface value to render. In practice, one usually biases the isosurface to a value that is below zero by a small fraction of the full dynamic range in order to avoid noise in regions where the velocity is close to zero.

### 3.2 Isosurface Extraction and Immersive Visualization

To develop an unsteady immersive rendering of the vortex evolution, we begin with a set of full-volume dumps of the primitive variables  $(u, v, w, p)$  at a selected number of time steps (typ. 100) covering one period of the vortex shedding cycle. For each time step, we compute  $\lambda_2$  at each grid point with the X-based postprocessor. The resultant scalar field is then interpolated onto an unstructured mesh along with the pressure. In both phases the high-order accuracy of the spectral element bases is essential to the construction of high-quality data sets and their subsequent projection onto meshes of arbitrary size. This data is then passed to a surface stripper (built on top of vtk) to extract isosurfaces at a given threshold  $\tau$  and produce a set of triangles corresponding to  $\lambda_2 = \tau$ . Typically,  $-2 \leq \tau \leq -1$ , out of a range of  $[-30, 40]$  for  $\lambda_2$ .

Generation of isosurfaces is a complex task. Because the extraction of a single isosurface from the volume data requires between 8–15 minutes per animation frame, offline processing of the surfaces is necessary. In addition to isosurface extraction, smoothing and triangulation algorithms are applied to the volume data at a negligible increase in run time. Finally, the surface files are saved to disk in the vtk polydata format. This format contains all the information necessary to construct the geometric representation of the surface with the corresponding pressure at each vertex point.

The surface files are then loaded into another vtk-based application that is built on top of the CAVE library (the vtk classes for renderer and renderwindow have been extended to operate within the CAVE [7]) and is designed to run on all members of the CAVE family of display technology, including the ImmersaDesk. The CAVE is a 10'x10'x10' cube of rear-projected screens, while the ImmersaDesk is single 5' rear-projected screen. Both are capable of displaying immersive stereo images. The immersive application makes full use of the CAVE library's rich set of tools to enable users to explore and interrogate individual surface files, as well as conduct the exploration over a series of surface files (typically a time series with fixed threshold or vice versa). For the  $Re = 700$  surface files, a rate of 15 frames per second was achieved. This was more than sufficient to produce a realistic sense of motion to the user. Control over viewing location, frame rate, and playback are provided via combination of the navigational wand and a virtual menu. In addition, users can define a color table to map other scalars onto the  $\lambda_2$  isosurfaces. At present, only the pressure is mapped.

Within the immersive environment it is also possible to generate high-quality renderings of the surface data. The user can navigate to points of interest within the data set and take virtual snapshots or virtual movies of the experience, capturing the needed information for high-resolution rendering. This is done either by immediately generating a RenderMan RIB file using a class provided within vtk for a snapshot, or by capturing the camera position and orientation information in virtual movie mode to allow off-line rendering. (Virtual movie mode requires an additional offline step for generating the RIB files, because the size of the files would interfere with the required real-time response of the virtual environment.) Once the RIB files are created, they are processed by using *Blue Moon Rendering Tools* (BMRT), a freely available collection of rendering tools that adhere to the RenderMan interface standard [8]. BMRT enables users to capture points of interest from within the data set in a high-quality format suitable for publication or videotape.




Figure 3: Symmetry plane data for  $Re = 700$ ,  $K = 1021$ ,  $N = 11$  (a-b), and  $N = 13$  (c-e): (a) mean velocity profiles,  $\langle u \rangle$ ; (b) rms velocity  $\langle u'^2 \rangle^{1/2}$ ; (c) pressure contours ; (d) spanwise vorticity,  $\omega_y$ ; (e) contours of  $\lambda_2 < 0$ .

## 4 Results

To put the dynamic results into perspective, we begin with a series of two-dimensional renderings in the symmetry plane,  $y = 0$ . Figure 3 shows the results of a simulation at  $Re = 700$  using 1021 elements of order  $N = 11$ . Time-averaged velocity profiles are shown in (a), and corresponding rms values,  $u_{rms} := \langle u'^2 \rangle^{1/2}$ , are shown in (b). Here  $u' := u - \langle u \rangle$  is the fluctuating component of  $u$ , and  $\langle . \rangle$  denotes a single-period time average. The profiles have maximum values of 1.075 and 0.282 for  $u$  and  $u_{rms}$ , respectively. We see in (a) the Blasius profile at the inlet, the separated and recirculating wake region at  $x = 2$ , followed by a gradual recovery until the outlet, where the mean velocity profile is fuller than the inflow profile.




Figure 4: Key vortex structures at  $Re = 700$ : (a) standing horseshoe vortex (b) interlaced tails (c) hairpin head (d) bridge. Contours mapped onto  $\lambda_2 = -1$  surface represent pressure (light=high, dark=low). Ceiling lighting has been added to enhance depth perspective and is responsible for the shadows cast on the plate.

From the rms profiles (b), it is clear that the flow upstream of the hemisphere is essentially steady. The immediate wake region at  $x = 2$  exhibits remarkably little unsteadiness. The passage of the hairpin vortices is evidenced by the strong rms fluctuations in the wake region, which also reveal the lifting of the vortices away from the plate. Further downstream, there is a significant growth in activity near the wall, as can be seen by the peak in the rms profile at the outlet.

To indicate the structure of individual hairpin vortices at a fixed instance, we present contours of several quantities in Fig. 3 (c–e). Pressure contours are shown in (c). The vortex cores are readily identified by the low-pressure zones, shown in bold. Contours of spanwise vorticity (d) also show the hairpin vortices and, in particular, the movement of the heads and tips away from the plate. Finally, (e) shows contours of  $\lambda_2 \in [-30, 0]$ , revealing the intersection of the vortex tips with the symmetry plane and the presence of a steady horseshoe vortex at the base of the hemisphere near  $x = -1$ .

While the views in Fig. 3 give a fair amount of information about the hairpin vortex evolution, they fail to reveal any three-dimensional details. Figures 4 and 5 show  $\lambda_2 = -1$  isosurfaces for  $Re = 700$ . We observe several vortex features, some of which we had not identified prior to viewing the unsteady animation in an immersive environment. Figure 4 shows the classic horseshoe vortex (a) upstream of the hemisphere, which is also commonly found upstream of end-mounted cylinders (as evidenced by snow drift patterns at the bases of trees and telephone poles). Moving downstream, we see the interlacing of the hairpin vortex tails (b), as observed by Acalar and Smith [1], the hairpin head (c), and a vortex “bridge” (d), which is a common form of vortex reconnection in viscous flows [10, 11, 14, 20]. Acalar and Smith [1] note that the bridge-head structure eventually separates from the hairpin and lifts off as a separate vortex ring. While we observe this phenomenon in our simulations, the ring is so quickly dissipated by viscosity at this Reynolds number that the liftoff is not pronounced.

Because the bridge is quite thin and rather short lived, significant resolution in space and time is






Figure 5: View showing distinction between “spikes” and tails of preceding vortex.

required to see it. We carefully chose the frame in order to present the bridge here. However, initial detection of the bridge and other similar unanticipated structures requires observation of a sequence of images from many viewing angles. Another structure detected as a result of such interactive viewing is shown in Fig. 5. The “spikes” seen jutting from the interior of the hairpin loop are readily visible in most of the higher Reynolds number computations we’ve done to date. They appear at just about the time that the tails of the preceding vortex disappear as a result of stretching-induced dissipation. Initially, we believed the spikes were remnants of the dissipating tails. However, careful observation of the animation revealed that the spikes and tails briefly appear at the same time, indicating that they are not part of the same vortex structure. Downstream views of the head-spike structure seem to indicate that the spike formation is induced by the close proximity of the tails to the head, as seen in Fig. 4b.

## 5 Discussion and Conclusion

We have examined the structure of hairpin vortices in the near wake region of a hemispherical roughness element at  $Re_R = 700$  using spectral element simulations coupled with interactive immersive visualization. This is part of a wider investigation, with Reynolds numbers ranging from 450 (just above transition) to 1200, into the role hairpin vortices play in the transition between laminar and turbulent flow. We have shown excellent agreement between simulation and experimental data with regard to shedding frequency and clear identification of the principal features of the primary hairpin vortex and its evolution. We found that the combination of immersion and motion played a key role in the assimilation, integration, and understanding of this data set and believe that immersive visualization tools, such as those developed here, will be essential tools for future investigation of flows where coherent structures play a key role in flow transition and development.

Further investigation is required into the nature of the secondary vortices that develop immediately

downstream of the recirculation zone. Acalar and Smith observed two types of secondary vortices. The first formed in the wake of the hairpin and was entirely contained between the legs. It is an open question as to whether the bridge we observed corresponds to this structure and whether this is an example of classic vortex reconnection via bridging [11]. The second is a pair of vortices that form several diameters downstream of the hemisphere on either side of the primary hairpin vortex. We do indeed see the formation of such structures, which become more pronounced in recent calculations at higher Reynolds numbers. Singer and Joslin have shown that a cascade of such hairpins ultimately evolves into a turbulent spot [17]. Their calculations were in a plane channel, which allowed them to track the (isolated) vortex as it moved downstream, resulting in considerable computational savings. A similar approach could be used here, exploiting the temporal periodicity of the solution to provide a well-defined inflow boundary condition. However, it is clear that straightforward calculation of spatially developing flows from transition to turbulence in inhomogeneous geometries will require a significant increase in resolution to capture the range of scales present in both the solution and geometry.

Exploration of extremely large data sets is difficult even with today’s state-of-the-art visualization tools. For example, to store the primitive variables  $(u, v, w, p)$  for a 150-frame “movie” requires 5–20 GB. Since we are interested in determining how vortex topology varies with Reynolds number, we anticipate having tens to hundreds of such “movies” in our database. There is a consequent need for advanced technology to archive, manipulate, and explore data sets in excess of a terabyte.

We are currently working on methods that facilitate such investigation. One promising path is to use subsampled or lower-resolution data sets for preliminary investigation and to record areas or paths where high-resolution images are desired. These images can then be computed offline at full data set resolutions. We are developing this and other visualization software technology, such as automatic feature detection and tracking, to assist in the interrogation of large data sets such as considered here.

## Acknowledgments

This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under Contract W-31-109-Eng-38, and by the Department of Energy under Grant No. B341495 to the Center on Astrophysical Thermonuclear Flashes at University of Chicago.

## References

- [1] M. S. Acalar and C. R. Smith, “A study of hairpin vortices in a laminar boundary layer: Part 1, hairpin vortices generated by a hemisphere protuberance”, *J. Fluid Mech.*, **175**, pp. 1–41 (1987).
- [2] D. C. Banks, T. Crockett, R. D. Joslin, B. A. Singer, “Parallel Rendering of Complex Vortical Flows”, <http://www.icase.edu/docs/hilites/banks/parallelRend.html>.

- [3] C. Cruz-Neira, D. J. Sandin and T. A. DeFanti, “Surround-screen projection-based virtual reality: The design and implementation of the CAVE<sup>TM</sup>”, *Computer Graphics* (Proceedings of SIGGRAPH ’93), ACM SIGGRAPH, August, pp. 135–142 (1993).
- [4] M. Dryja and O. B. Widlund, “An additive variant of the Schwarz alternating method for the case of many subregions”, *Technical Rep. 339*, Dept. Comp. Sci., Courant Inst., NYU (1987).
- [5] P. F. Fischer, “An overlapping Schwarz method for spectral element solution of the incompressible Navier-Stokes equations”, *J. Comp. Phys.*, **133**, pp. 84–101 (1997).
- [6] P. F. Fischer, N. I. Miller, and H. M. Tufo, “An overlapping Schwarz method for spectral element simulation of three-dimensional incompressible flows”, in *Parallel Solution of Partial Differential Equations*, P. Bjorstad and M. Luskin, eds., Springer-Verlag, New York, (1999).
- [7] Futures Lab Web page, [www.mcs.anl.gov/fl](http://www.mcs.anl.gov/fl).
- [8] L. Gritz and J. K. Hahn, “BMRT: A global illumination implementation of the RenderMan standard”, *J. Graphics Tools*, **1**(3), pp. 29–47 (1996).
- [9] J. Jeong and F. Hussain, “On the identification of a vortex”, *J. Fluid Mech.*, **285**, pp. 69–94 (1995).
- [10] S. Kida, M. Takaoka, and F. Hussain, “Reconnection of two vortex rings”, *Phys. Fluids A*, **1**(4), pp. 630–632 (1989).
- [11] S. Kida and M. Takaoka, “Vortex reconnection”, *Annu. Rev. Fluid Mech.*, **26**, pp. 169–189 (1994).
- [12] P. S. Klebanoff, W. G. Cleveland, and K. D. Tidstrom, “On the evolution of a turbulent boundary layer induced by a three-dimensional roughness element”, *J. Fluid Mech.*, **92**, pp. 101–187 (1992).
- [13] Y. Maday and A. T. Patera, “Spectral element methods for the Navier-Stokes equations”, in *State of the Art Surveys in Computational Mechanics*, A. K. Noor, ed., ASME, New York, pp. 71–143 (1989).
- [14] M. V. Melander and F. Hussain, “Cross-linking of two antiparallel vortex tubes” *Phys. Fluids A*, **1**(4), pp. 633–636 (1989).
- [15] W. Schroeder, K. Martin, and B. Lorensen, The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics, Prentice Hall, Englewood Cliffs, N. J. (1998).
- [16] B. A. Singer “The formation and growth of a hairpin vortex”, *Instability, Transition, and Turbulence*, M. Y. Hussaini, A. Kumar, and C. L. Street, eds., Springer-Verlag, New York, pp.367 (1992).
- [17] B. A. Singer and R. D. Joslin, “Metamorphosis of a Hairpin Vortex Into a Young Turbulent Spot”, *Physics of Fluids*, **6** (11), pp. 3724–3736 (1994).
- [18] C. R. Smith, J. D. A. Walker, A. H. Haidari, and U. Sobrun, “On the dynamics of near-wall turbulence”, *Phil. Trans. R. Soc. Lon. A*, **336**, pp. 131–175 (1991).

- [19] H. M. Tufu and P. F. Fischer, “Fast parallel direct solvers for coarse grid problems”, *J. Parallel and Distributed Computing* (to appear).
- [20] N. J. Zabusky and M. V. Melander, “Three-dimensional vortex tube reconnection: morphology for orthogonally-offset tubes”, *Physica D*, **37**, pp. 555–562 (1989).